DOE Climate Change Prediction Program

Warren Washington and Jerry Meehl, National Center for Atmospheric Research; Bert Semtner, Naval Postgraduate School; John Weatherly, U.S. Army Cold Regions Research and Engineering Lab Laboratory.

During FY02, a substantial resource was allocated at NERSC to the DOE Climate Change Prediction Program to complete a 1000-year control simulation of the present climate with the new Community Climate System Model (CCSM2) developed at the National Center for Atmospheric Research (NCAR). CCSM2 is a state-of-the-art model which tightly couples individual, independently developed, atmosphere, land, ocean, and sea ice components of the earth's climate system. The previous version of the model, released in 1996, was one of the first fully coupled models which produced a stable simulation of surface temperature without the need for flux correction (ad hoc adjustments to the energy exchange between the atmosphere and ocean). CCSM2 is an enhanced version of the first model, with improved physical process parameterizations and upgraded computational methodologies for better parallel performance on clustered machines such as the IBM SP at NERSC.

A 1000-year simulation demonstrates the ability of CCSM2 to produce a long-term, stable representation of the earth's climate. Few if any climate models in the world can make this claim, since all previous simulations contained drifts too large to allow complete, uncorrected simulations to 1000 years. In addition, the simulation provides scientists with a database to analyze the variability of weather and climate on time scales ranging from interannual to interdecadal to intercentennial. Few datasets exist which are as comprehensive as the one produced during this simulation.

Computationally, the full CCSM2 code is complex, and consists of five binaries which are organized to execute concurrently within a single job. The models exchange data at various frequencies appropriate to the physical, large-scale processes being simulated. The standard message passing library MPI is used to facilitate the exchange of information. CCSM2 requires 4.5 wall-clock hours on 144 1.5-Gflops CPUs of the NERSC IBM SP to complete one simulated year. Over 90% of the 760,000 processor hours allocated on the IBM SP have been used as of July 2002. NERSC gave CCSM2 special queue priority to complete this project in a timely fashion. Results of this simulation have not yet been published, but preliminary results of 800 model years were presented to 250 participants at the Seventh Annual CCSM Workshop held in Breckenridge, Colorado, on June 25–27, 2002.

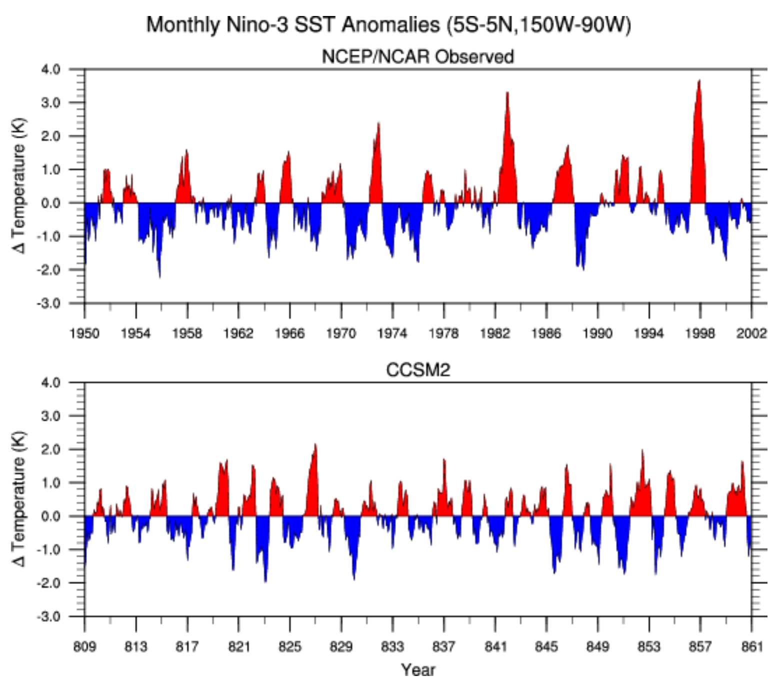

A time series of monthly averaged sea surface temperature anomalies in the central Pacific Ocean (the Niño-3 region), which provide a measure of the strength and frequency of the periodic El Niño event. The upper panel shows observations during the 1950–2002 period, and the lower panel shows years 809-861 from the CCSM2 simulation. CCSM2 simulates the magnitude of the El Niño phenomenon relatively well, but with too frequent an occurrence.

High-Resolution Global Coupled Ocean/Sea Ice Modeling

Mathew E. Maltrud, Los Alamos National Laboratory; Julie L. McClean, Naval Postgraduate School.

The objective of this project is to couple a high-resolution ocean general circulation model with a high-resolution dynamic-thermodynamic sea ice model in a global context. Currently, such simulations are typically performed with a horizontal grid resolution of about 1 degree. At this resolution (about 30 to 50 km in the polar regions), the ocean model cannot resolve very narrow current systems (including fronts and turbulent eddies) that play a crucial role in the transport of heat and salt in the global ocean. Similarly, lower-resolution sea ice models cannot resolve important dynamics that occur in regions of complicated topography (such as the Canadian Archipelago).

This project is running a global ocean circulation model with horizontal resolution of approximately 1/10th degree (between 11 km and 2.5 km). This is the highest-resolution simulation even attempted with a such a realistic model. This configuration has dimensions of $3600 \times 2400 \times 40$, resulting in 177 million active ocean grid points (some grid points are on land). The code being used is the Parallel Ocean Program (POP), developed at LANL under the Department of Energy's CHAMMP program. At NERSC, 448 processors are used to run the model. One year can be simulated in about eight wall-clock days (86,000 processor hours), generating over 500 GB of output. Eight model years have been run to date, with a goal of 30–50 years. After the ocean simulation has run for 10–15 model years, it will be coupled with a sea ice model to more accurately simulate the polar circulation. Approximately 385,000 processor hours have been used of the 920,000 hours allocated for this project. No new results or publications have been produced yet, since the model is still equilibrating.

The interaction of the ocean and overlying sea ice in global coupled numerical models is poorly understood, though very important. When ocean water freezes into sea ice, salt is released into the upper ocean, making it more dense. Conversely, when the ice melts, it creates a layer of fresh water that is less dense than the underlying ocean. This delicate balance between melting and freezing is very difficult to simulate with coarse grids. In particular, high vertical resolution is needed near the surface to simulate this salinity balance correctly. High horizontal resolution is required to properly simulate the current systems that advect these salinity anomalies into the open ocean. Inaccuracies in the surface ocean properties due to poor representation of ocean-ice interaction can have wide-ranging global consequences. Most notable is the possibility that too much fresh surface water can inhibit vertical convection in the northern seas (since it is less dense than the salty water beneath it), which then disrupts the entire global heat budget. Coarse-resolution simulations have found that the circulation and heat budget are extremely sensitive to the way sea ice is prescribed in ocean-only runs. The best tool for simulating the global circulation accurately is a high-resolution, fully coupled ocean-sea ice model.

1/10 Degree Global POP Ocean Model Currents at 50m Depth
(blue = 0; red > 150 cm/s)

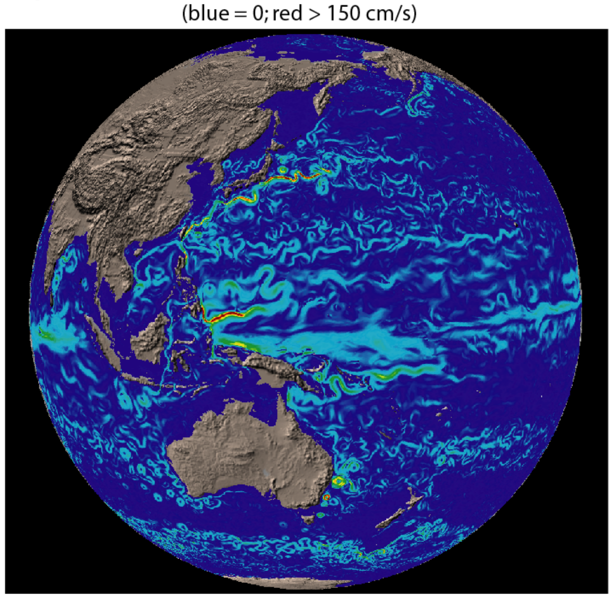

Supernova Explosions and Cosmology

Peter Nugent and Daniel Kasen, Lawrence Berkeley National Laboratory; Peter Hauschildt, University of Georgia; Edward Baron, University of Oklahoma; Stan Woosley and Gary Glatzmaier, University of California, Santa Cruz; Tom Clune, Goddard Space Flight Center; Adam Burrows, Salim Hariri, Phil Pinto, Hessam Sarjoughian, and Bernard Ziegler, University of Arizona; Chris Fryer and Mike Warren, Los Alamos National Laboratory; Frank Dietrich and Rob Hoffman, Lawrence Livermore National Laboratory.

This collaboration brings together the SciDAC Supernova Science Center and the members of the PHOENIX/SYNPOL collaboration. The goal is a better understanding of supernovae of all types through simulation and model validation. Specific objectives are to clarify the physics of supernova explosions, to improve the reliability of such explosions as calibrated standard candles, and to measure fundamental cosmological parameters. Despite decades of research and modeling, no one understands in detail how supernovae work. The problem persists largely because, until recently, computer resources have been inadequate to carry out credible multi-dimensional calculations.

On June 4, 2002, at the American Astronomical Society meeting in Albuquerque, N.M., Michael Warren and Chris Fryer from Los Alamos National Laboratory presented the results of one of several projects in this collaboration, the first 3-D supernova explosion simulation, based on computation at NERSC. This research eliminates some of the doubts about earlier 2-D modeling and paves the way for rapid advances on other questions about supernovae.

Earlier one-dimensional simulations of core-collapse supernovae almost always failed to explode. Two-dimensional simulations were qualitatively different from 1-D, leading to a robust explosion without fine-tuning of the star's physical properties. They showed that the explosion process is critically dependent on convection, the mixing of the matter surrounding the iron core of the collapsing star. It was believed that the results could again be changed radically by adding a third dimension, but the 3-D simulations turned out to be similar to the 2-D results. The explosion energy, explosion time scale, and remnant neutron star mass do not differ by more than 10 percent between the 2-D and 3-D models. With these 3-D results, researchers are ready to attack more exotic problems that involve rotation and non-symmetric accretion.

The 3-D simulation used a parallel smooth particle hydrodynamics (SPH) code coupled with a flux-limited diffusion radiation transport. Supernova calculations are computationally demanding because many processes, involving all four fundamental forces of physics, must be modeled and followed for more than 100,000 time steps. Typical simulations (1 million particles) took about three months on the IBM SP at NERSC. This collaboration has used approximately 660,000 of its allocated 1 million hours.

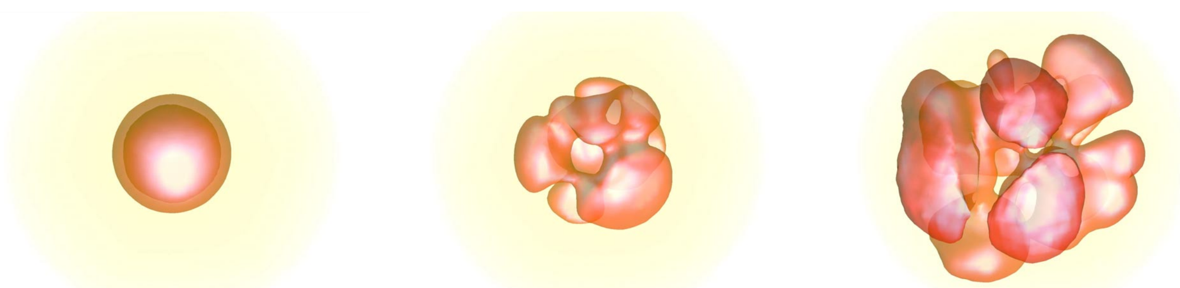

Computer visualization shows (left to right) three stages of a simulated supernova explosion over a period of 50 milliseconds, starting about 400 milliseconds after the core begins to collapse. The surfaces show the material which is flowing outward at a speed of 1000 kilometers/second. Left is the initial spherical implosion. Center, as in-falling gas approaches the core, it is exposed to a higher and higher influx of neutrinos that heat the gas and make it buoyant. Right, as more cold gas sinks in, it is heated and rises, resulting in enough convective energy transfer to create an explosion. (Michael S. Warren, Los Alamos National Laboratory)

Chris L. Fryer and Michael S. Warren, "Modeling Core-Collapse Supernovae in 3-Dimensions," Astrophysics Journal (in press).

Hamuy, M. & Pinto, P. A., "Type II Supernovae as Standardized Candles," Astrophysical Journal **566,** L63 (2002).

Fryer, C. L., Holz, D. E., & Hughes, S. A., "Gravitational Wave Emission from Core Collapse of Massive Stars," Astrophysical Journal **565,** 430 (2002).

Black Hole Merger Simulations

Ed Seidel, Gabrielle Allen, Denis Pollney, and Peter Diener, Max Planck Institute for Astrophysics; John Shalf, Lawrence Berkeley National Laboratory.

This group is performing simulations of the spiraling coalescence of two black holes, a problem of particular importance for interpreting the gravitational wave signatures that will soon be seen by new laser interferometric detectors around the world. Detection of the first gravitational waves (or failure to do so) will strongly test Einstein's Theory of General Relativity, the results of which will have ramifications that extend throughout the world of physics. The Cactus simulation code is being used to perform the calculations. These simulations must use half of the NERSC IBM SP's available aggregate memory of 4.2 TB in order to achieve the resolution required to accurately simulate these phenomena. This is the first time ever that a spiraling merger of this type has been accurately simulated.

Collisions between black holes should theoretically create propagating gravitational waves, similar to the electromagnetic waves given off by distant stars. These ripples in space-time should be seen as subtle variations in the length of objects as they move through space. Recently built laser interferometric detectors such as LIGO and VIRGO are capable of measuring these subtle ripples in space. However, the gravitational wave signal that can be detected by these interferometers is so faint that it is very close to the level of noise in these devices. So simulations of the kinds of events that might produce gravitational waves can provide important insights into the gravitational wave signature produced by these events, potentially making the instruments more productive.

The Cactus code performs a direct evolution of Einstein's equations, which are a system of coupled nonlinear elliptic hyperbolic equations that contain millions of terms if fully expanded. Consequently, the simulation resource requirements are enormous just to do the most basic of simulations. The simulations have been limited by both the memory and CPU performance of supercomputers as they attempt to move from calibrating against analytic black hole solutions to non-analytic astrophysically relevant cases in full 3-D. The spiraling merger is just such a non-analytic case.

This simulation uses 1.5 terabytes of memory and more than 2 terabytes of disk space for each run on the NERSC IBM SP system. These runs typically consume 64 of the large-memory nodes of the SP (a total of 1024 processors) for 48 wall-clock hours at a stretch. The simulation can use all 184 nodes, but this would only allow simulations that are fractionally larger than using the large-memory nodes due to memory/load-balancing issues. In the space of two months, these simulations consumed 700,000 of the allocated 760,000 CPU hours, simulating three-fourths of a full orbit before coalescence. The results so far indicate that the Meudon model for coalescence seems to match the simulation data more accurately than the competing Cook-Baumgarte model.

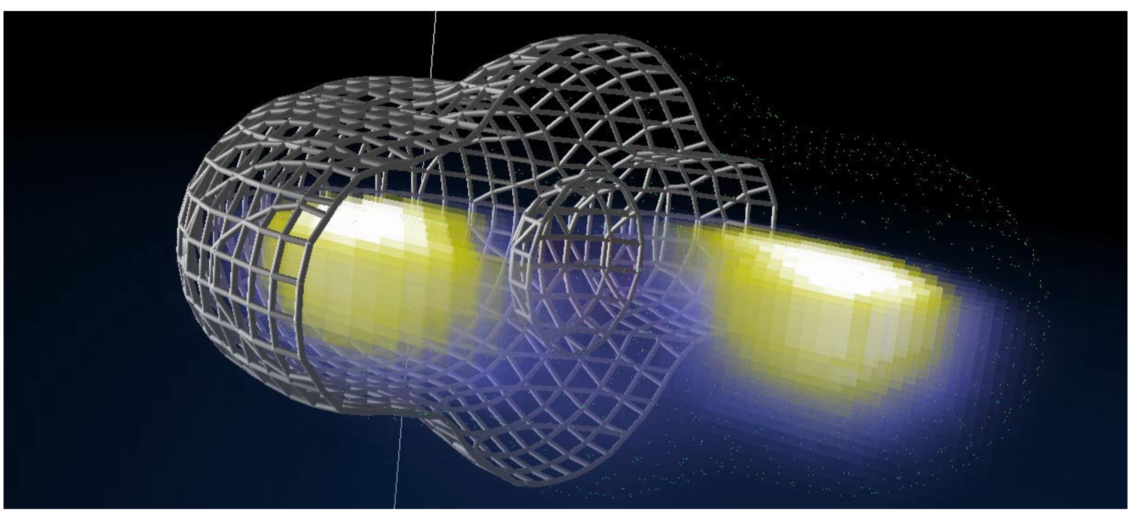

Results related to this work, including visualizations of binary black hole inspiral, appeared in the April 2002 edition of Scientific American, on the Discovery Channel in June 2002, in the June 2002 IEEE Computer Magazine and will also appear in Nature in the near future.